Exhibit 14

CHART FOR U.S. PATENT NO. 8,452,929 ("the '929 Patent")

Accused Products

Dell's products, including but not limited to Dell's PowerScale All-flash (*e.g.*, F910, F710, F210. F900, F600, F200 Isilon F800, and Isilon F810), Hybrid (*e.g.*, H7000, and H700), and Archive (*e.g.*, A3000, and A300) products with PowerScale OneFS functionality ("Accused Products"), infringe at least Claim 35 of the '929 Patent.

Claims	Exemplary Evidence of Infringement
[35 pre] A system for storage of data, the system comprising:	To the extent the preamble is limiting, the Accused Products comprise a system for storage of data. For example, the Accused Products "comprise[ ] scale-out file storage platforms configured with the OneFS operating system," where "OneFS combines three layers of traditional storage architectures . . . into one unified software layer, creating a single intelligent distributed file system that runs on a OneFS powered storage cluster." For example, "the PowerScale enterprise-class storage platform includes all-flash, hybrid and archive nodes" and the "software defined architecture of PowerScale OneFS . . . powers the world's most secure NAS storage array." *See, e.g.*: The PowerScale family comprises of scale-out file storage platforms configured with the OneFS operating system. OneFS provides the intelligence behind the highly scalable, high–performance modular storage solution that can grow with your business. A PowerScale OneFS cluster can be built with a flexible choice of storage platforms including all-flash, hybrid and archive nodes. These solutions provide performance, choice, efficiency, flexibility, scalability, security, and protection for you to store massive amounts of unstructured data within a cluster. The PowerScale all-flash nodes co-exist seamlessly in the same cluster with your existing PowerScale or Isilon nodes to drive your traditional workloads and even the most modern applications like Generative AI. The PowerScale all-flash storage platforms include:

Claims	Exemplary Evidence of Infringement
	PowerScale F910  **PowerScale F910** is the latest in our next-generation all-flash nodes lineup and provides massive AI-ready performance with the ultimate capacity in a highly dense 2U configuration. Each node hosts 24 NVMe SSDs. F910 allows you to scale raw storage from 92 TB to 737 TB per node and up to 186 PB of raw capacity per cluster. The F910 includes in-line compression and deduplication to maximize efficiency (Energy Star certification coming soon). The minimum number of PowerScale nodes per cluster is three while the maximum cluster size is 252 nodes. The F910 is best suited for high-capacity workloads within demanding verticals like media and entertainment, high frequency trading, healthcare and accelerating phases of the AI lifecycle for Generative AI applications. PowerScale F710  Our next-generation **PowerScale F710**, leveraging PowerEdge R660, delivers high performance and improved density in a 1U platform with up to 10 all-flash NVMe SSD drives per node. The F710 allows you to scale raw storage from 38 TB to 307 TB per node and up to 77 PB of raw capacity per cluster. The F710 includes in-line compression and deduplication, with Energy Star certification coming soon. The minimum number of PowerScale nodes per cluster is three while the maximum cluster size is 252 nodes. The F710 is best suited for Generative AI and AI workloads, as well as high performing vertical workloads like, media and entertainment, healthcare and life sciences, high frequency trading, and EDA workloads. PowerScale F210  **PowerScale F210** is also part of our next-generation all-NVMe lineup. It delivers significant performance gains over the previous generation in a cost-effective 1U form factor with up to 4 NVME all-flash SSD drives per node. The F210 offers a 15TB QLC option and allows you to scale raw storage from 8 TB to 61 TB per node and up to 15 PB of raw capacity per cluster. It also includes in-line compression and deduplication, with Energy Star certification coming soon. The minimum number of PowerScale nodes per cluster is three while the maximum cluster size is 252 nodes. The F210 is best suited for customer beginning their AI and Analytics journey, and other high-demanding workloads that require a balance of performance and capacity. Dell PowerScale All-Flash

Claims	Exemplary Evidence of Infringement
	OneFS combines the three layers of traditional storage architectures—file system, volume manager, and data protection—into one unified software layer, creating a single intelligent distributed file system that runs on a OneFS powered storage cluster. 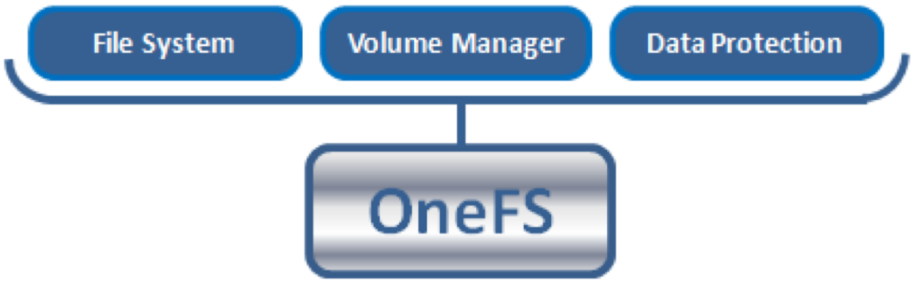 **Figure 1. OneFS combines file system, volume manager, and data protection into one single intelligent, distributed system** Dell PowerScale OneFS: Technical Overview White Paper

Claims	Exemplary Evidence of Infringement

PowerScale nodes

The following are the supported PowerScale nodes:

Table 24. PowerScale nodes

Node	OneFS Compatibility									
	9.0.x	9.1.x	9.2.0.x	9.2.1.x	9.3.x	9.4.x	9.5.x	9.6.x	9.7.x	9.8.x
A3000				✓	✓	✓	✓	N/A	✓	✓
A300				✓	✓	✓	✓	N/A	✓	✓
H7000				✓	✓	✓	✓	N/A	✓	✓
H700				✓	✓	✓	✓	N/A	✓	✓
F900			✓	✓	✓	✓	✓	N/A	✓	✓
F600	✓	✓	✓	✓	✓	✓	✓	N/A	✓	✓
F200	✓	✓	✓	✓	✓	✓	✓	N/A	✓	✓
B100					✓	✓	✓	N/A	✓	✓
P100					✓	✓	✓	N/A	✓	✓
F210								N/A	✓	✓
F710								N/A	✓	✓
F910								N/A		✓

PowerScale OneFS Supportability and Compatibility Guide

Designed to handle your most ambitious data challenges, the PowerScale enterprise-class storage platform includes all-flash, hybrid and archive nodes as well as multicloud solutions. The software defined architecture of PowerScale OneFS, the operating system that powers the world's most secure NAS storage array[2], enables simplicity at scale, intelligent insights, and the ability to place the data anywhere it needs to be – at the edge, in the core or in the cloud. PowerScale can be deployed and consumed anywhere your data is – in your on-premises data center as an appliance, in multi-cloud and native cloud solutions or in APEX delivered as-a-service. Whether you are hosting file shares, home directories or delivering high performance data access for applications like analytics, AI/ML, video rendering or life sciences, PowerScale can seamlessly scale performance, capacity, and efficiency to handle any unstructured data workload to drive both, traditional and modern applications.

Dell Powerscale Data Sheet

<table>
<tr><th>Claims</th><th>Exemplary Evidence of Infringement</th></tr>
<tr><td>[35 a] a volatile memory device;</td><td>
The Accused Products comprise a volatile memory device.
For example, the Accused Products “take advantage of DDR5” and include “Dual Rank DDR5 RDIMMs.” For example, the Accused Products’ “caching infrastructure leverages the system memory (RAM),” where “L1 Cache” and “L2 Cache” are “RAM” with “Persistence’ characterized as “Volatile” and “hold” or “contain[] . . . file system data and metadata blocks . . . .” For example, “level 1 (L1) and level 2 (L2)[] are memory (RAM) based and analogous to the cache used in processors (CPUs).”
See, e.g.:
The F210 and F710 offer greater density in a 1U platform, with the F710 supporting 10 NVMe SSDs per node and the F210 offering a 15.36 TB drive option. The Sapphire Rapids CPU provide 19% lower cycles-per-instruction. PCIe Gen 5 doubles throughput when compared to PCIe Gen 4. Further, the nodes take advantage of DDR5, offering greater speed and bandwidth.
Table 3. PowerScale F710 node specifications
<table>
<tr><th>Attribute</th><th>PowerScale F710 Specification</th></tr>
<tr><td>Chassis</td><td>1U Dell PowerEdge R660</td></tr>
<tr><td>CPU</td><td>Dual Socket – Intel Sapphire Rapids 6442Y (2.6G/24C)</td></tr>
<tr><td>Memory</td><td>Dual Rank DDR5 RDIMMs 512 GB (16 x 32 GB)</td></tr>
<tr><td>Journal</td><td>1 x 32 GB SDPM</td></tr>
<tr><td>Front-end networking</td><td>2 x 100 GbE or 25 GbE</td></tr>
<tr><td>Infrastructure networking</td><td>2 x 100 GbE</td></tr>
<tr><td>NVMe SSD drives</td><td>10</td></tr>
</table>
</td></tr>
</table>

<table>
<tr><th>Claims</th><th>Exemplary Evidence of Infringement</th></tr>
<tr><td></td><td>

Table 1. PowerScale F210 node specifications

Attribute	PowerScale F210 Specification
Chassis	1U Dell PowerEdge R660
CPU	Single Socket – Intel Sapphire Rapids 4410Y (2G/12C)
Memory	Dual Rank DDR5 RDIMMs 128 GB (8 x 16 GB)
Journal	1 x 32 GB SDPM
Front-end networking	2 x 100 GbE or 25 GbE
Infrastructure networking	2 x 100 GbE or 25 GbE
NVMe SSD drives	4

PowerScale All-Flash F210 and F710 AI Ready Data Platform White Paper

The OneFS caching infrastructure leverages the system memory (RAM) and nonvolatile memory (NVRAM) in each node. It uses solid state drives (SSDs) and the increased capacity, affordability, and storage persistence that they offer.

Dell PowerScale OneFS SmartFlash File System Caching Infrastructure

</td></tr>
</table>

Claims	Exemplary Evidence of Infringement
	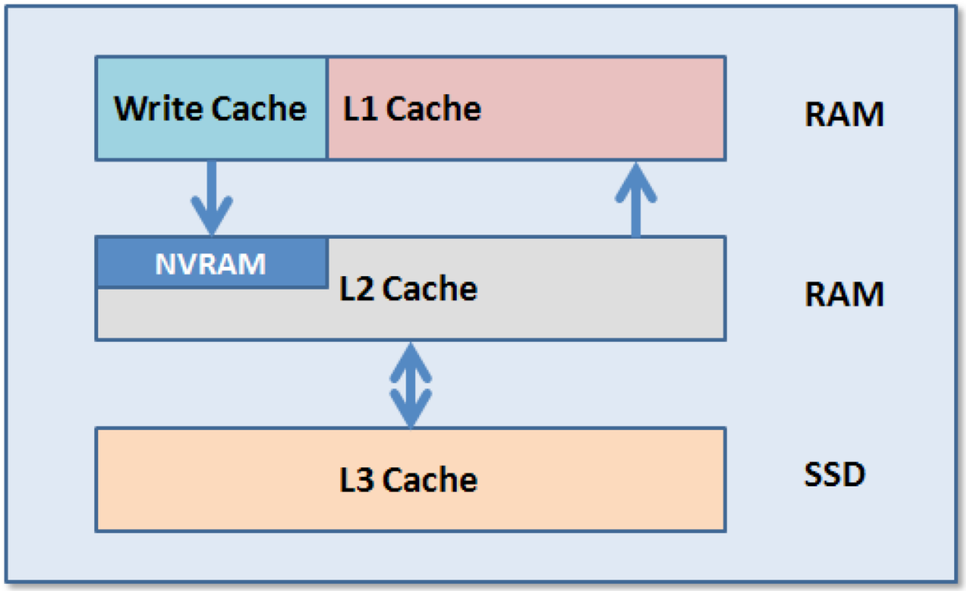 **Figure 9. OneFS caching hierarchy** The first two types of read cache, level 1 (L1) and level 2 (L2), are memory (RAM) based and analogous to the cache used in processors (CPUs). These two cache layers are present in all platform storage nodes.

<table>
<tr><th>Claims</th><th>Exemplary Evidence of Infringement</th></tr>
<tr><td></td><td>
Table 2. OneFS cache
<table>
<tr><th>Name</th><th>Type</th><th>Persistence</th><th>Description</th></tr>
<tr><td>L1 Cache</td><td>RAM</td><td>Volatile</td><td>Also called front-end cache, holds clean, cluster coherent copies of file system data and metadata blocks requested by clients over the front-end network</td></tr>
<tr><td>L2 Cache</td><td>RAM</td><td>Volatile</td><td>Back-end cache, containing clean copies of file system data and metadata on a local node</td></tr>
<tr><td>SmartCache / Write Coalescer</td><td>NVRAM</td><td>Non-volatile</td><td>Persistent, battery backed NVRAM journal cache which buffers any pending writes to front-end files that have not been committed to disk.</td></tr>
<tr><td>SmartFlash L3 Cache</td><td>SSD</td><td>Non-volatile</td><td>Contains file data and metadata blocks evicted from L2 cache, effectively increasing L2 cache capacity.</td></tr>
</table>
Dell PowerScale OneFS: Technical Overview
</td></tr>
<tr><td>[35 b] a non-volatile memory; and</td><td>
The Accused Products comprise a non-volatile memory.

For example, the Accused Products' "caching infrastructure leverages nonvolatile memory (NVRAM)," using "solid state drives and the increased capacity, affordability, and storage persistence they offer." For example, the Accused Products' "SmartCache/Write Coalescer" are "NVRAM" with "Non-volatile" "Persistence" described as "journal cache." For example, the Accused Products' "SmartFlash L3 Cache" are "SSD[s]" with "Non-Volatile" "Persistence" that "contain[] file data and metadata blocks." For example, the Accused Products include "Drive Media" that are "Flash NVMe/SAS," "SATA/SAS & SSD," and/or "SATA."

See, e.g.:
</td></tr>
</table>

Claims	Exemplary Evidence of Infringement
	The OneFS caching infrastructure leverages the system memory (RAM) and nonvolatile memory (NVRAM) in each node. It uses solid state drives (SSDs) and the increased capacity, affordability, and storage persistence that they offer. Dell PowerScale OneFS SmartFlash File System Caching Infrastructure 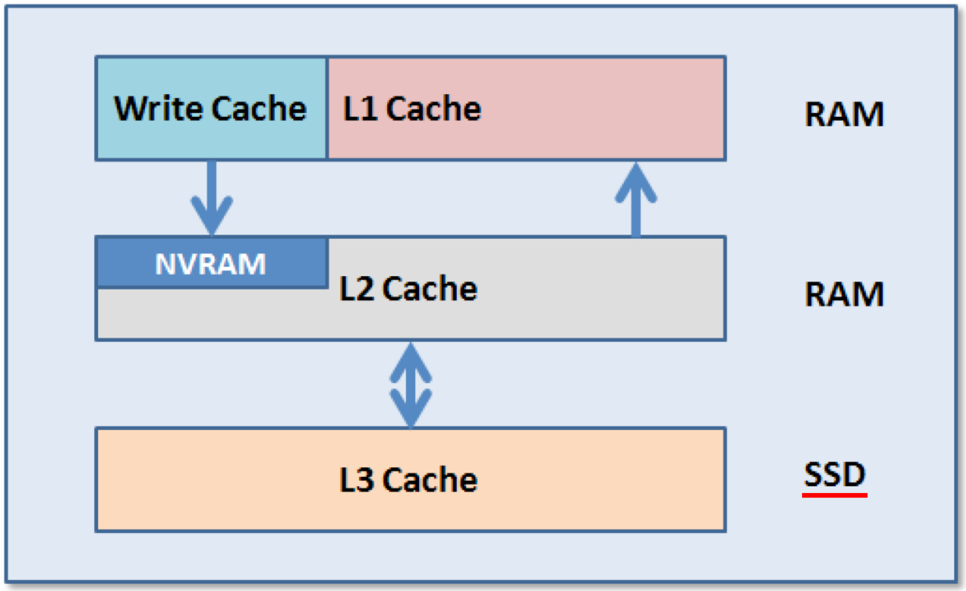

Claims	Exemplary Evidence of Infringement
	Table 2. OneFS cache

Name	Type	Persistence	Description
L1 Cache	RAM	Volatile	Also called front-end cache, holds clean, cluster coherent copies of file system data and metadata blocks requested by clients over the front-end network
L2 Cache	RAM	Volatile	Back-end cache, containing clean copies of file system data and metadata on a local node
SmartCache / Write Coalescer	NVRAM	Non-volatile	Persistent, battery backed NVRAM journal cache which buffers any pending writes to front-end files that have not been committed to disk.
SmartFlash L3 Cache	SSD	Non-volatile	Contains file data and metadata blocks evicted from L2 cache, effectively increasing L2 cache capacity.

<table>
<tr><th>Claims</th><th>Exemplary Evidence of Infringement</th></tr>
<tr><td></td><td>
The OneFS powered platform nodes are broken into several classes, or tiers, according to their functionality:
<table>
<tr><th>Tier</th><th>I/O Profile</th><th>Drive Media</th><th>Nodes</th></tr>
<tr><td>Performance</td><td>High Perf, Low Latency</td><td>Flash NVMe/SAS</td><td>F900 F810
F600 F800
F200</td></tr>
<tr><td>Hybrid / Utility</td><td>Concurrency & Streaming Throughput</td><td>SATA/SAS & SSD</td><td>H700 H600
H7000 H5600
H500
H400</td></tr>
<tr><td>Archive</td><td>Nearline & Deep Archive</td><td>SATA</td><td>A300 A200
A3000 A2000</td></tr>
</table>
Figure 2. Hardware tiers and nodes: Types
Dell PowerScale OneFS: Technical Overview
</td></tr>
<tr><td>[35 c] a processor in communication with the volatile memory device and the non-volatile memory device, wherein the</td><td>
The Accused Products include a processor in communication with the volatile memory device and the non-volatile memory device.

For example, the Accused Products include “Intel Xeon CPUs”/an “Intel Xeon Processor E5-2697A v4”/a “Dual Socket Intel Processor”/a “Single Socket Intel Processor” that are “configured with the OneFS operating system,” where the “software defined architecture of PowerScale OneFS” combines “three layers of traditional storage architectures . . . into one unified software layer . . . that runs on a OneFS powered storage cluster.”

See, e.g.:
</td></tr>
</table>

Claims	Exemplary Evidence of Infringement
processor is configured to	**AI-ready performance**. Leveraging continuous innovation of the OneFS software and the latest technology (i.e., Intel Xeon CPUs, the latest DRAM PCIe), the F910 delivers faster time to AI insights with up to 127% improved streaming performance.[4] It accelerates the model checkpointing an training phases of the AI pipeline and keeps GPUs utilized with up to 300 PBs of storage per cluster. PowerScale also recently introduced greate flexibility and option to access data in the cloud with the launch of Dell APEX File Storage for Azure, delivering 6x greater cluster performance co to Azure NetApp Files.[5] [Dell PowerScale: AI Success Starts With AI-Ready Storage]() The PowerScale family comprises of scale-out file storage platforms configured with the OneFS operating system. OneFS prov the intelligence behind the highly scalable, high–performance modular storage solution that can grow with your business. A PowerScale OneFS cluster can be built with a flexible choice of storage platforms including all-flash, hybrid and archive nodes. These solutions provide performance, choice, efficiency, flexibility, scalability, security, and protection for you to store massive amounts of unstructured data within a cluster. The PowerScale all-flash nodes co-exist seamlessly in the same cluster with your existing PowerScale or Isilon nodes to drive traditional workloads and even the most modern applications like Generative AI. The PowerScale all-flash storage platforms inc [Dell PowerScale All-Flash]() **Isilon F800 All-Flash Specifications**

F800 ATTRIBUTES & OPTIONS	1.6 TB SSD	3.2 TB SSD	3.84 TB SSD	7.68 TB SSD
Raw chassis capacity	96 TB	192 TB	230 TB	460 TB
SSD drives (2.5") per chassis			60	
Self-Encrypting drive (SED SSD) FIPS 140-2 compliant option			Yes	
Operating system		OneFS 8.1 or later except for self-encrypting drive options which 8.1.0.1 or later		
Number of nodes per chassis			4	
CPU type (per node)			Intel® Xeon® Processor E5-2697A v4	

Claims	Exemplary Evidence of Infringement
	PowerScale F600 All-Flash Specifications F600 ATTRIBUTES & OPTIONS | 1.92 TB SSD | 3.84 TB SSD Raw node capacity | 15.36 TB | 30.72 TB NVMe SSD drives (2.5") per node | 8 Self-encrypting drive (SSD SED) option | No Operating system | PowerScale OneFS 9.0 CPU type (per node) | Dual Socket Intel® Processor **PowerScale F200 All-Flash Specifications** F200 ATTRIBUTES & OPTIONS | 960 GB SSD | 1.92 TB SSD Raw node capacity | 3.84 TB | 7.68 TB SSD drives (2.5") per node | 4 Self-encrypting drive (SED SSD) option | No Operating system | PowerScale OneFS 9.0 CPU type (per node) | Single socket Intel® Processor Specification Sheet Dell EMC PowerScale Designed to handle your most ambitious data challenges, the PowerScale enterprise-class storage platform includes all-flash, hybrid and archive nodes as well as multicloud solutions. The software defined architecture of PowerScale OneFS, the operating system that powers the world's most secure NAS storage array[2], enables simplicity at scale, intelligent insights, and the ability to place the data anywhere it needs to be – at the edge, in the core or in the cloud. PowerScale can be deployed and consumed anywhere your data is – in your on-premises data center as an appliance, in multi-cloud and native cloud solutions or in APEX delivered as-a-service. Whether you are hosting file shares, home directories or delivering high performance data access for applications like analytics, AI/ML, video rendering or life sciences, PowerScale can seamlessly scale performance, capacity, and efficiency to handle any unstructured data workload to drive both, traditional and modern applications.

Claims	Exemplary Evidence of Infringement
	Dell Powerscale Data Sheet OneFS combines the three layers of traditional storage architectures—file system, volume manager, and data protection—into one unified software layer, creating a single intelligent distributed file system that runs on a OneFS powered storage cluster.  **Figure 1. OneFS combines file system, volume manager, and data protection into one single intelligent, distributed system** Dell PowerScale OneFS: Technical Overview White Paper
[35 c][i][1] maintain a dynamic table in volatile storage representing data stored in the non-volatile memory,	The Accused Products maintain a dynamic table in volatile storage representing data stored in the non-volatile memory. For example, the Accused Products include a "LIN (logical inode) table" that "maps a LIN to blocks on the disk for the inode" and "maps the LIN, SnapID that identifies a file version to the location on disk that houses the metadata about that file." For example, the Accused Products "creat[e] snapshot[s]" where a "snapshot is basically a logical pointer to data that is stored on a cluster at a particular point in time." For example, "creating a snapshot involves synchronizing changes across several levels, from the write-back cache (coalescer) to the snapshot tracking file (STF)," including "generat[ing]" a "snapshot . . . up to the snapshot root" and "creat[ing] and updat[ing]" a "snapshot tracking file." *See, e.g.*:

Claims	Exemplary Evidence of Infringement
	A OneFS snapshot is basically a logical pointer to data that is stored on a cluster at a particular point in time. Each snapshot references a specific directory under OneFS and includes all the files stored in that directory and its subdirectories. If the data referenced by a snapshot is modified, the snapshot stores a physical copy of the data that was modified. Snapshots are either created according to user configuration or are automatically generated by OneFS to facilitate system operations. LIN table Finding the next version of a particular file is a fast operation due to the structure of the LIN (logical inode) table. While the LIN table is frequently called a table, it is actually a B-Tree. It is sorted by (LIN,Version), so finding the next newer version of a LIN is an inexpensive operation. A LIN is fundamentally the identity of an object and possess the following attributes: • Directories point to LINs. • The LIN Table maps a LIN to blocks on disk for the inode. • The Inode contains the map of data ranges to blocks on disk. • The LIN table maps the LIN,SnapID that identifies a file version to the location on disk that houses the metadata about that file. Given a LIN:SnapID pairing, finding the next higher SnapID is a fast and efficient process, and the HEAD version is always represented by the biggest SnapID. On the backend, creating a snapshot involves synchronizing changes across several levels, from the write-back cache (coalescer) to the snapshot tracking file (STF). The main elements of this process are: 1. The snapshot acquires an exclusive snapshot lock. 2. Coalescers are coordinated and suspended across all nodes. 3. The write lock is upgraded to exclusive. 4. The maximum snapshot ID is incremented. 5. The snapshot (and associated minisnaps) is generated up to the snapshot root. 6. A snapshot tracking file is created and updated. 7. If requested, a snapshot alias is created. Data Protection with Dell PowerScale SnapshotIQ
[35 c][i][2] a checkpointed version of the dynamic table in the non-	The Accused Products maintain a checkpointed version of the dynamic table in the non-volatile memory. For example, the Accused Products "record" any "changes to the dataset . . . in the pertinent snapshot inodes, which contain only referral ('ditto') records until any of the logical blocks they reference are altered or another snapshot is taken" and "the first time a file is changed after a snapshot is taken, the LIN is recorded in the snapshot tracking file." For example, when "a file's metadata is changed," the Accused Products "need[] to create a copy of

Claims	Exemplary Evidence of Infringement
volatile memory, and	the data, and increment the Snap_ID" and "the first time a file is changed after a snapshot is taken, the LIN is recorded in the snapshot tracking file (STF)." For example, the Accused Products' "OneFS Snapshots . . . only [store] the changed blocks of a file . . . when updating the snapshots." For example, a "snapshot will initially be written to the default data pool and then moved." *See, e.g.*: The process of capturing a snapshot in OneFS is near instantaneous. However, there is a small amount of snapshot preparation w that has to occur. The moment a snapshot is taken, it consumes zero space until any file writes occur to the data it contains. Any changes to a dataset are then recorded in the pertinent snapshot inodes, which contain only referral ('ditto') records, until any c the logical blocks they reference are altered or another snapshot is taken. In order to reconstruct data from a particular snapshot, OneFS will look though all of the more recent versions snapshot tracking files (STFs) until it reaches HEAD (current version). In so doing, it will systematically find all the changes and recreate the point-in-time view of that dataset. **Snapshot tracking files** Snapshot tracking files (STF) are the main data structure associated with a snapshot. A snapshot tracking file has three major purposes: • Indicating which snapshots are active. • Storing snapshot attributes, such as usage data, creation time, and root directory paths. • Recording a list of LINs modified in the snapshot, which can be freed when the snapshot is deleted. Snapshot Tracking Files are a special file type with several unique characteristics, and are involved in the full snapshot life cycle, including the creation, storing any changes, and deletion of snapshots.

Claims	Exemplary Evidence of Infringement
	Writing to a snapshot This following figure illustrates the process of updating a file's metadata, by way of a simple UID change. When this occurs, and OneFS detects that it needs to duplicate the old data, it copies the old blocks to the snapshot. In this case, there was no LIN 1:abcd:1234/98 prior to the write, so any attempt to read this LIN would have fallen forward to the head version. 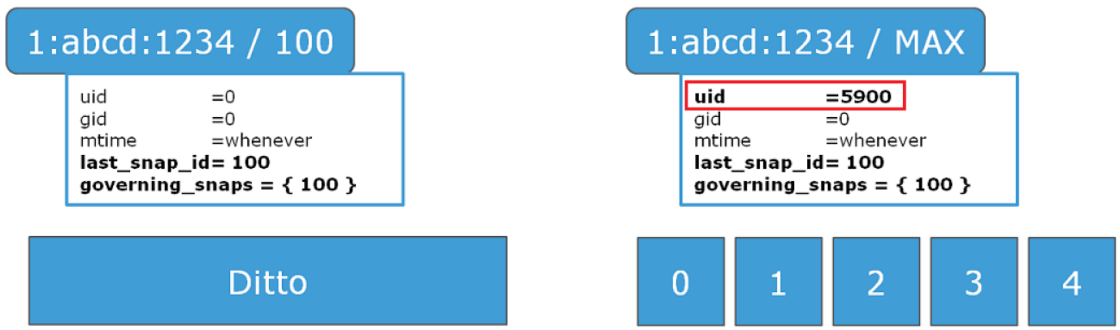 **Figure 10: Changing the UID of a File in a Snapshot – Part 1** When a file's metadata is changed, OneFS needs to create a copy of the data, and increment the Snap_ID. Since there is no actual data written, there is no need to copy over any old data blocks to the new snapshot, so ditto-blocks are used instead. However, the first time a file is changed after a snapshot is taken, the LIN is recorded in the snapshot tracking file (STF). This STF lets OneFS efficiently know what data can be removed when a snapshot is deleted. For example, the snapshots taken on a performance aligned tier can be physically housed on a more cost-effective archive tier. The snapshot storage target setting is applied to each file version by SmartPools. When a snapshot is taken, the storage pool setting is simply preserved. This means that the snapshot will initially be written to the default data pool and then moved. The SmartPools job subsequently finds the snapshot version and moves it to the intended pool during the next scheduled SmartPools job run. Data Protection with Dell PowerScale SnapshotIQ

<table>
<tr><th>Claims</th><th>Exemplary Evidence of Infringement</th></tr>
<tr><td></td><td>OneFS Snapshots are highly scalable and typically take less than one second to create. They create little performance overhead, regardless of the level of activity of the file system, the size of the file system, or the size of the directory being copied. Also, only the changed blocks of a file are stored when updating the snapshots, thereby ensuring highly efficient snapshot storage utilization. User access to the available snapshots is through a /.snapshot hidden directory under each file system directory.

High Availability and Data Protection with Dell PowerScale Scale-Out NAS</td></tr>
<tr><td>[35 c][i][3] a log representing transactions affecting stored data, and</td><td>The Accused Products maintain a log representing transactions affecting stored data.

For example, the Accused Products use “journals for transaction logging” where the “file system journal . . . stores information about changes to the file system” and “journal[s] all the transactions that are occurring across every node in the storage cluster,” “ensuring that writes [] are safe across power failures.”

See, e.g.:

For write operations, where coherency is vital, the BAM first sets up a transaction. Next it uses a 2-phase commit protocol (2PC) over the RBM to guarantee the success of an atomic write operation across all participant nodes. This is managed using the BAM Safe Write (BSW) code path. The 2PC atomically updates multiple disks across the 2PC participant nodes, using their journals for transaction logging. The write path operates as

Each node or chassis added to a cluster increases aggregate disk, cache, CPU, and network capacity. OneFS leverages each of the hardware building blocks, so that the whole becomes greater than the sum of the parts. The RAM is grouped together into a single coherent cache, allowing I/O on any part of the cluster to benefit from data cached anywhere. A file system journal ensures that writes that are safe across power failures. Spindles and CPU are combined to increase throughput, capacity and IOPS as the cluster grows, for access to one file or for multiple files. A cluster’s storage capacity can range from tens of TBs up to 186 PB of raw capacity. The maximum capacity will continue to increase as storage media and node chassis continue to get denser.

High Availability and Data Protection with Dell PowerScale Scale-Out NAS</td></tr>
</table>

Claims	Exemplary Evidence of Infringement
	Every node that owns blocks in a particular write is involved in a two-phase commit. The mechanism relies on NVRAM for journaling all the transactions that are occurring across every node in the storage cluster. Using multiple NVRAMs in parallel allows for high-throughput writes while maintaining data safety against all manner of failures, including power failures. In the event that a node should fail mid-transaction, the transaction is restarted instantly without that node involved. When the node returns, the only required actions are for the node to replay its journal from NVRAM—which takes seconds or minutes—and, occasionally, for AutoBalance to rebalance files that were involved in the transaction. No expensive 'fsck' or 'disk-check' processes are ever required. No drawn-out resynchronization ever needs to take place. Writes are never blocked due to a failure. This patented transaction system is one of the ways that OneFS eliminates single—and even multiple—points of failure.

Claims	Exemplary Evidence of Infringement
	The initiator top half of the "captain" node uses a modified two-phase commit transaction to safely distribute writes to multiple NVRAMs across the cluster, as shown in the following figure. 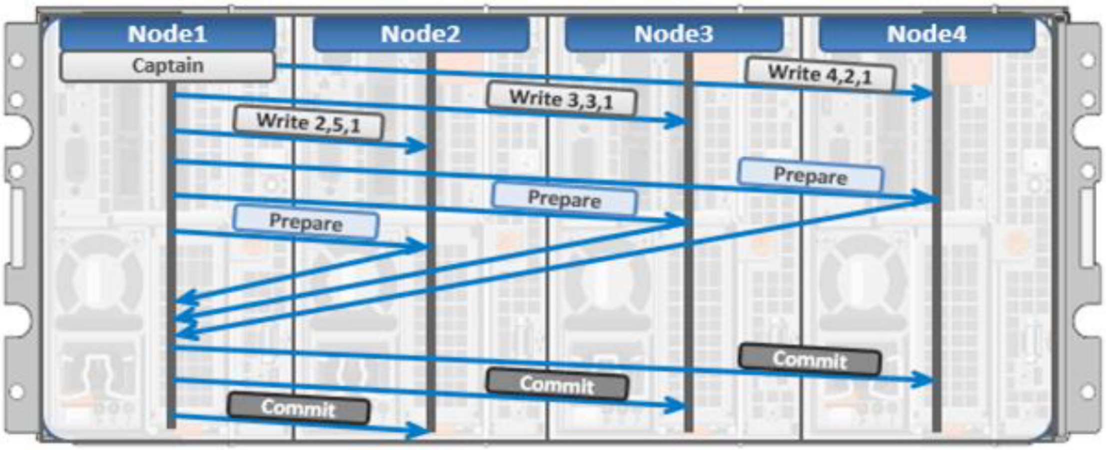 **Figure 8. Distributed transactions and two-phase commit** Every node that owns blocks in a particular write is involved in a two-phase commit. The mechanism relies on NVRAM for journaling all the transactions that are occurring across every node in the storage cluster. Using multiple NVRAMs in parallel allows for high-throughput writes while maintaining data safety against all manner of failures, including power failures. In the event that a node should fail mid-transaction, the transaction is restarted instantly without that node involved. When the node returns, the only required actions are for the node to replay its journal from NVRAM—which takes seconds or minutes—and, occasionally, for AutoBalance to rebalance files that were involved in the transaction. No expensive 'fsck' or 'disk-check' processes are ever required. No drawn-out resynchronization ever needs to take place. Writes are never blocked due to a failure. This patented transaction system is one of the ways that OneFS eliminates single—and even multiple—points of failure.

Claims	Exemplary Evidence of Infringement
	A file system journal, which stores information about changes to the file system, is designed to enable fast, consistent recoveries after system failures or crashes, such as power loss. The file system replays the journal entries after a node or cluster recovers from a power loss or other outage. Without a journal, a file system would need to examine and review every potential change individually after a failure (an “fsck” or “chkdsk” operation); in a large file system, this operation can take a long time. OneFS is a journaled file system in which each node contains a battery-backed NVRAM card used for protecting uncommitted writes to the file system. The NVRAM card battery charge lasts many days without requiring a recharge. When a node boots up, it checks its journal and selectively replays transactions to disk where the journaling system deems it necessary. Dell PowerScale OneFS: Technical Overview **Writing files** On a node, the input-output operations of the OneFS software stack split into two functional layers: A top layer, or initiator, and a bottom layer, or participant. In read and write operations, the initiator and the participant play different roles. When a client writes a file to a node, the initiator on the node manages the layout of the file on the cluster. First, the initiator divides the file into blocks of 8 KB each. Second, the initiator places the blocks in one or more stripe units. At 128 KB, a stripe unit consists of 16 blocks. Third, the initiator spreads the stripe units across the cluster until they span a width of the cluster, creating a stripe. The width of the stripe depends on the number of nodes and the protection setting. After dividing a file into stripe units, the initiator writes the data first to non-volatile random-access memory (NVRAM) and then to disk. NVRAM retains the information when the power is off. During the write transaction, NVRAM guards against failed nodes with journaling. If a node fails mid-transaction, the transaction restarts without the failed node. When the node returns, it replays the journal from NVRAM to finish the transaction. The node also runs the AutoBalance job to check the file's on-disk striping. Meanwhile, uncommitted writes waiting in the cache are protected with mirroring. As a result, OneFS eliminates multiple points of failure. **The file system journal** A journal, which records file-system changes in a battery-backed NVRAM card, recovers the file system after failures, such as a power loss. When a node restarts, the journal replays file transactions to restore the file system. PowerScale OneFS Web Administration Guide 9.3.0.0

Claims	Exemplary Evidence of Infringement
	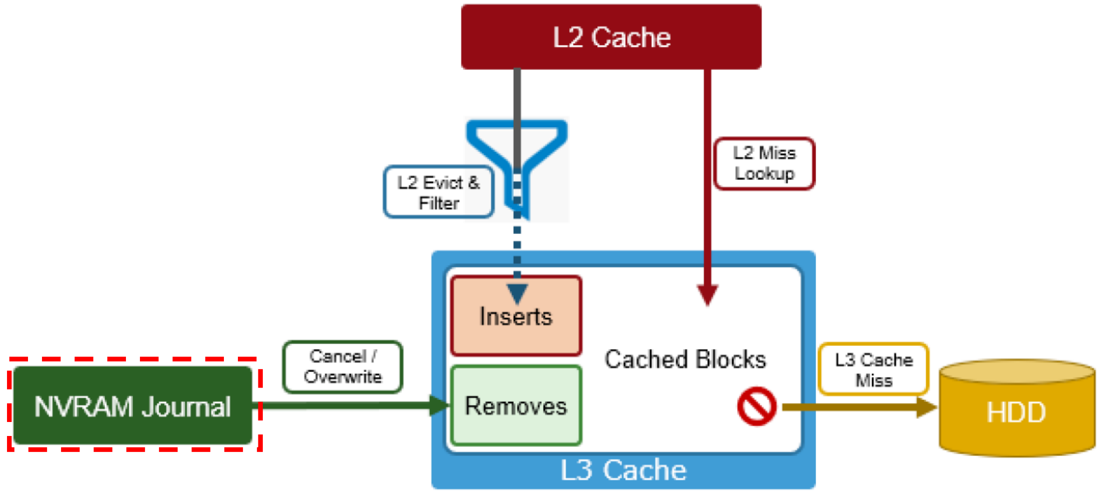 Dell PowerScale OneFS SmartFlash File System Caching Infrastructure
[35 c][ii] to repair the dynamic table by discovering data written to the non-volatile memory between the last update of the log and a time of repair.	The Accused Products repair the dynamic table by discovering data written to the non-volatile memory between the last update of the log and a time of repair. For example, “in order to reconstruct data from a particular snapshot,” the Accused Products “look through all of the more recent versions of the snapshot tracking files (STFs) until it reaches the HEAD (current version). In so doing, it will systematically find all of the changes and recreate the point-in-time view of that dataset.” For example, the Accused Products are “a journaled file system . . . [and w]hen a node boots up, it checks its journal and selectively replays transactions to a disk where the journaling system deems it necessary.” For example, in “the event that” an Accused Product “should fail mid-transaction, . . . the only required actions are for the node to replay its journal . . . .” For example, the Accused Products’ “FlexProtect locates any unprotected files on the cluster and repairs them as quickly as possible,” and includes “Drive Scan,” “LIN verification,” and “Device Removal” phases. For example, the Accused Products’ “IntegrityScan . . . is responsible for examining the entire file system for inconsistencies . . . [i]f IntegrityScan detects a checksum mismatch, a system alert is generated and written to the syslog and OneFS automatically attempts to repair the suspect block.”

Claims	Exemplary Evidence of Infringement
	See, e.g.: OneFS is a journaled file system in which each node contains a battery-backed NVRAM card used for protecting uncommitted writes to the file system. The NVRAM card battery charge lasts many days without requiring a recharge. When a node boots up, it checks its journal and selectively replays transactions to disk where the journaling system deems it necessary. Every node that owns blocks in a particular write is involved in a two-phase commit. The mechanism relies on NVRAM for journaling all the transactions that are occurring across every node in the storage cluster. Using multiple NVRAMs in parallel allows for high-throughput writes while maintaining data safety against all manner of failures, including power failures. In the event that a node should fail mid-transaction, the transaction is restarted instantly without that node involved. When the node returns, the only required actions are for the node to replay its journal from NVRAM—which takes seconds or minutes—and, occasionally, for AutoBalance to rebalance files that were involved in the transaction. No expensive 'fsck' or 'disk-check' processes are ever required. No drawn-out resynchronization ever needs to take place. Writes are never blocked due to a failure. Dell PowerScale OneFS: Technical Overview The process of capturing a snapshot in OneFS is near instantaneous. However, there is a small amount of snapshot preparation w that has to occur. The moment a snapshot is taken, it consumes zero space until any file writes occur to the data it contains. Any changes to a dataset are then recorded in the pertinent snapshot inodes, which contain only referral ('ditto') records, until any the logical blocks they reference are altered or another snapshot is taken. In order to reconstruct data from a particular snapshot, OneFS will look though all of the more recent versions snapshot tracking files (STFs) until it reaches HEAD (current version). In so doing, it will systematically find all the changes and recreate the point-in-time view of that dataset. Data Protection with Dell PowerScale SnapshotIQ

Claims	Exemplary Evidence of Infringement
	Run automatically after a drive or node removal or failure, FlexProtect locates any unprotected files on the cluster and repairs them as quickly as possible. The FlexProtect job includes the following distinct phases: • **Drive Scan.** FlexProtect scans the cluster's drives, looking for files and inodes in need of repair. When such file or inode is found, the job opens the LIN and repairs it and the corresponding data blocks using the restripe process. • **LIN Verification.** Once the drive scan is complete, the LIN verification phase scans the inode (LIN) tree and verifies, reverifies, and resolves any outstanding reprotection tasks. • **Device Removal.** In this final phase, FlexProtect removes successfully repaired drives or nodes from the cluster. Certain jobs such as FlexProtect have a corresponding job provided with a name suffixed by `Lin`, for example FlexProtectLin. This indicates that the job will automatically use an SSD-based copy of metadata, where available, to scan the LIN tree, rather than the drives themselves. Depending on the workflow, this often significantly improves job runtime performance. OneFS marks blocks that are actually in use by the file system. IntegrityScan, for example, traverses the live file system, marking every block of every LIN in the cluster to proactively detect and resolve any issues with the structure of data in a cluster. The jobs in the marking exclusion set are: • Collect • IntegrityScan • MultiScan Dell PowerScale OneFS Job Engine White Paper

Claims	Exemplary Evidence of Infringement
	IntegrityScan, another component of the OneFS Job Engine, is responsible for examining the entire file system for inconsistencies. It does this by systematically reading every block and verifying its associated checksum. Unlike traditional fsck-style file system integrity checking tools, IntegrityScan is designed to run while the cluster is fully operational, removing the need for any downtime. If IntegrityScan detects a checksum mismatch, a system alert is generated and written to the syslog and OneFS automatically attempts to repair the suspect block. High Availability and Data Protection with Dell PowerScale Scale-Out NAS